Claims

What is claimed is:

1. A system for maintaining a high availability processing environment, said system comprising:

a network having a plurality of clusters, each cluster of the network comprising a plurality of identical servers, each cluster of the network being directly connected to at least one other cluster of the network, wherein each pair of clusters directly connected to each other is characterized by each server in a first cluster of the pair of clusters being directly connected to at least one server in a second cluster of the pair of clusters via a communication link; and

a control server adapted to monitor an operational status of said communication link, said operational status of the communication link being that said communication link is operational or non-operational, said control server being directly linked to at least one server in each cluster via a communication channel between the control server and the at least one server.

2. The system of claim 1, further comprising:

a global dataset that includes an identification of each communication link in the network, said global dataset being accessible to the control server; and

a local dataset specific to each cluster of the plurality of clusters, said local dataset including an identification of each communication link in the network to which the servers of said each cluster is coupled for flow of data out of the cluster, said global dataset being

accessible to the servers of said each cluster.

3. The system of claim 1, wherein the control server is adapted to monitor an operational status of a first communication link between a first server of the first cluster and a second server of the second cluster by sending a query signal to the first server, said query signal requesting the first server to send a response signal to the control server indicating the operational status of the first communication link, said operational status of the first communication link being that said first communication link is operational or non-operational.

4. The system of claim 3, wherein the first server is adapted to respond to the query signal by sending a prompt signal over the first communication link to the second server, said prompt signal prompting the second server to send a return signal to the first server over the first communication link, said return signal or absence thereof being indicative of the operational status of the first communication link.

5. The system of claim 3, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational, the control server is adapted to notify the load balancer that the first communication link is non-operational.

6. The system of claim 3, wherein upon being notified that the first communication link is non-operational, the load balancer is adapted to fail over the first server with respect to the first communication link.

7. The system of claim 3, wherein upon receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational, the control server is adapted to inform a service node in the first cluster that the first communication link is non-operational.

8. The system of claim 7, wherein upon being informed that the first communication link is non-operational, the service node is adapted to make a determination of a cause of the first communication link being non-operational.

9. The system of claim 8, wherein upon making said determination of said cause, the service node is adapted to facilitate making the first communication link operational.

10. The system of claim 3, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server, the control server is adapted make a determination that the first server is non-operational and to notify the load balancer that the first server is non-

operational.

11. The system of claim 10, wherein upon being notified that the first server is non-operational, the load balancer is adapted to fail over the first server.

12. The system of claim 10, wherein upon not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server, the control server is adapted to inform a service node in the first cluster that the first server is non-operational.

13. The system of claim 12, wherein upon being informed that the first server is non-operational, the service node is adapted to make a determination of a cause of the first server being non-operational.

14. The system of claim 13, wherein upon making said determination of said cause the service node is adapted to facilitate making the first server operational.

15. The system of claim 1, wherein each cluster of the plurality of clusters has a load balancer adapted to distribute data traffic uniformly among the servers comprised by each cluster.

16. The system of claim 1, wherein the control server is adapted to receive a message from a first

server of the first cluster or from a load balancer of the first cluster, said message indicating that an entity is non-operational, said entity being selected from the group consisting of a server of the first cluster and a communication link between the first server of the first cluster and a second server of the second cluster.

17. The system of claim 16, wherein upon receiving said message the control server is adapted to inform a service node of the first cluster that the entity is non-operational.

18. The system of claim 1, wherein the control server is directly linked to a first server of the first cluster and is not directly linked to a second server of the first cluster, wherein the first server is directly connected to the second server, and wherein the control server is adapted to monitor an operational status of the second server via direct communication with the first server coupled with direct communication between the first server and the second server, said operational status of the second server being that said second server is operational or non-operational.

19. The system of claim 1, wherein at least one cluster of the plurality of clusters does not have a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster.

20. The system of claim 1, wherein the plurality of clusters includes a web cluster of web servers, an application cluster of application servers, and a database cluster of database servers, the web

cluster being directly connected to the application cluster, the application cluster being directly connected to the database cluster, the web cluster adapted to communicate with the database cluster by way of the application cluster functioning as an intermediary cluster between the web cluster and the database cluster.

21. The system of claim 20, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

22. The system of claim 20, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

23. The system of claim 1, wherein the plurality of clusters includes a web cluster of web servers and a database cluster of database servers, the web cluster being directly connected to the

database cluster, the web cluster adapted to directly communicate with the database cluster.

24. The system of claim 23, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

25. The system of claim 23, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

26. A method for maintaining a high availability processing environment, said method comprising:

providing a network having a plurality of clusters, each cluster of the network comprising a plurality of identical servers, each cluster of the network being directly connected to at least one other cluster of the network, wherein each pair of clusters directly connected to each other is characterized by each server in a first cluster of the pair of clusters being directly connected to at least one server in a second cluster of the pair of clusters via a communication link; and

providing a control server adapted to monitor an operational status of said communication link, said operational status of the communication link being that said communication link is operational or non-operational, said control server being directly linked to at least one server in each cluster via a communication channel between the control server and the at least one server.

27. The method of claim 26, said method further comprising:

providing a global dataset that includes an identification of each communication link in the network, said global dataset being accessible to the control server; and

providing a local dataset specific to each cluster of the plurality of clusters, said local dataset including an identification of each communication link in the network to which the servers of said each cluster is coupled for flow of data out of the cluster, said global dataset being accessible to the servers of said each cluster.

28. The method of claim 26, said method further comprising monitoring an operational status of

a first communication link between a first server of the first cluster and a second server of the second cluster, said monitoring being performed by the control server, said monitoring including sending a query signal to the first server, said query signal requesting the first server to send a response signal to the control server indicating the status of the first communication link, said operational status of the first communication link being that said first communication link is operational or non-operational.

29. The method of claim 28, wherein the first server is adapted to respond to the query signal by sending a prompt signal over the first communication link to the second server, said prompt signal prompting the second server to send a return signal to the first server over the first communication link, said return signal or absence thereof being indicative of the operational status of the first communication link.

30. The method of claim 28, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon the control server receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational the method further comprises: notifying the load balancer that the first communication link is non-operational, said notifying being performed by the control server.

31. The method of claim 28, wherein upon the load balancer being notified that the first

communication link is non-operational the method further comprises: failing over the first server with respect to the first communication link, said failing over being performed by the load balancer.

32. The method of claim 28, wherein upon the control server receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational the method further comprises: informing a service node in the first cluster that the first communication link is non-operational, said informing being performed by the control server.

33. The method of claim 32, wherein upon the service node being informed that the first communication link is non-operational the method further comprises: making a determination of a cause of the first communication link being non-operational, said making a determination being performed by the service node.

34. The method of claim 33, wherein upon the service node making said determination of said cause the method further comprises: facilitating making the first communication link operational, said facilitating being performed by the service node.

35. The method of claim 28, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon the

control server not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server the method further comprises: making a determination that the first server is non-operational and notifying the load balancer that the first server is non-operational, said making said determination that the first server is non-operational and said notifying being performed by the control server.

36. The method of claim 35, wherein upon load balancer being notified that the first server is non-operational the method further comprises: failing over the first server, said failing over being performed by the load balancer.

37. The method of claim 35, wherein upon the control server not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server the method further comprises: informing a service node in the first cluster that the first server is non-operational, said informing being performed by the control server.

38. The method of claim 37, wherein upon the service node being informed that the first server is non-operational the method further comprises: making a determination of a cause of the first server being non-operational, said making said determination of said cause being performed by the service node.

39. The method of claim 38, wherein upon the service node making said determination of said

cause the method further comprises: facilitating making the first server operational, said facilitating being performed by the service node.

40. The method of claim 26, wherein each cluster of the plurality of clusters has a load balancer, and wherein the method further comprises distributing data traffic uniformly among the servers comprised by each cluster, said distributing being performed by the load balancer.

41. The method of claim 26, wherein the control server is adapted to receive a message from a first server of the first cluster or from a load balancer of the first cluster, said message indicating that an entity is non-operational, said entity being selected from the group consisting of a server of the first cluster and a communication link between the first server of the first cluster and a second server of the second cluster.

42. The method of claim 41, wherein upon the control server receiving the message the method further comprises: informing a service node of the first cluster that the entity is non-operational, said informing being performed by the control server.

43. The method of claim 26, wherein the control server is directly linked to a first server of the first cluster and is not directly linked to a second server of the first cluster, wherein the first server is directly connected to the second server, said method further comprising monitoring the operational status of the second server via direct communication with the first server coupled

with direct communication between the first server and the second server, said monitoring being performed by the control server, said operational status of the second server being that said second server is operational or non-operational.

44. The method of claim 26, wherein at least one cluster of the plurality of clusters does not have a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster.

45. The method of claim 26, wherein the plurality of clusters includes a web cluster of web servers, an application cluster of application servers, and a database cluster of database servers, the web cluster being directly connected to the application cluster, the application cluster being directly connected to the database cluster, the web cluster adapted to communicate with the database cluster by way of the application cluster functioning as an intermediary cluster between the web cluster and the database cluster.

46. The method of claim 45, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

47. The method of claim 45, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

48. The method of claim 26, wherein the plurality of clusters includes a web cluster of web servers and a database cluster of database servers, the web cluster being directly connected to the database cluster, the web cluster adapted to directly communicate with the database cluster.

49. The method of claim 48, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

50. The method of claim 48, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.